

# An Efficient Approach for Computing Optimal Low-Rank Regularized Inverse Matrices

Julianne Chung

Virginia Tech

Department of Mathematics

Blacksburg, VA, USA

jmchung@vt.edu

Matthias Chung

Virginia Tech

Department of Mathematics

Blacksburg, VA, USA

mcchung@vt.edu

April 8, 2014

## Abstract

Standard regularization methods that are used to compute solutions to ill-posed inverse problems require knowledge of the forward model. In many real-life applications, the forward model is not known, but training data is readily available. In this paper, we develop a new framework that uses training data, as a substitute for knowledge of the forward model, to compute an optimal low-rank regularized inverse matrix directly, allowing for very fast computation of a regularized solution. We consider a statistical framework based on Bayes and empirical Bayes risk minimization to analyze theoretical properties of the problem. We propose an efficient rank update approach for computing an optimal low-rank regularized inverse matrix for various error measures. Numerical experiments demonstrate the benefits and potential applications of our approach to problems in signal and image processing.

*Keywords:* ill-posed inverse problems, regularization, low-rank approximation, TSVD, Bayes risk, empirical Bayes risk, machine learning

## 1 Introduction

Inverse problems arise in scientific applications such as biomedical imaging, geophysics, computational biology, computer graphics, and security [1, 2, 3, 4, 5, 6, 7, 8]. Assume that  $\mathbf{A} \in \mathbb{R}^{m \times n}$  and  $\mathbf{b} \in \mathbb{R}^m$  are given, then the linear *inverse problem* can be written as

$$\mathbf{A}\boldsymbol{\xi} + \boldsymbol{\delta} = \mathbf{b}, \tag{1}$$

where  $\boldsymbol{\xi} \in \mathbb{R}^n$  is the desired solution and  $\boldsymbol{\delta} \in \mathbb{R}^m$  is additive noise. We assume the underlying problem is *ill-posed*. A problem is ill-posed if the solution does not exist, is not unique, or does not depend continuously on the data [9]. A main challenge of ill-posed inverse problems is that small errors in the data may result in large errors in the reconstruction. In order to obtain a more meaningful reconstruction, regularization in the form of prior knowledge is needed to stabilize the inversion process.

Standard regularization methods require knowledge of  $\mathbf{A}$  [10, 2]. However, in many applications, the operator  $\mathbf{A}$  may not be readily available, or is only partially known. In this work, we investigate an approach that avoids including  $\mathbf{A}$  in the problem formulation, and instead learns the necessary information from training or calibration data. Assume training data  $\mathbf{b}^{(k)}, \boldsymbol{\xi}^{(k)}$  for  $k = 1, \dots, K$ , for (1) are given. First, our goal is to find a rank- $r$  matrix  $\mathbf{Z} \in \mathbb{R}^{n \times m}$  that gives a small error for the given training set. That is, the *sample mean error*,

$$f_K(\mathbf{Z}) = \frac{1}{K} \sum_{k=1}^K e_k(\mathbf{Z}) = \frac{1}{K} \sum_{k=1}^K \rho(\mathbf{Z}\mathbf{b}^{(k)} - \boldsymbol{\xi}^{(k)}) \quad (2)$$

should be small for some given error measure  $\rho : \mathbb{R}^n \rightarrow \mathbb{R}_0^+$ , e.g.,  $\rho(\mathbf{x}) = \|\mathbf{x}\|_p^p$ . Here,  $e_k(\mathbf{Z}) = \rho(\mathbf{Z}\mathbf{b}^{(k)} - \boldsymbol{\xi}^{(k)})$  refers to the *sample error* for a single observation  $k$ . In this paper, we consider efficient methods for obtaining an *optimal low-rank regularized inverse matrix*,

$$\hat{\mathbf{Z}} = \arg \min_{\text{rank}(\mathbf{Z}) \leq r} \frac{1}{K} \sum_{k=1}^K \rho(\mathbf{Z}\mathbf{b}^{(k)} - \boldsymbol{\xi}^{(k)}) . \quad (3)$$

Then, once matrix  $\hat{\mathbf{Z}}$  is computed, solving the inverse problem (1) requires a simple matrix-vector multiplication:  $\boldsymbol{\xi} = \hat{\mathbf{Z}}\mathbf{b}$ . Matrix  $\hat{\mathbf{Z}}$  can be precomputed, so that regularized solutions to inverse problems can be obtained efficiently and accurately.

Methods that incorporate training data have been widely used in machine learning [11, 12] and have been analyzed in the statistical learning literature [13]. Most previous work on using training data for solving inverse problems has been restricted to minimizing the predictive mean squared error and require knowledge of the forward model [14, 15]. In this paper, we provide a novel framework to solve inverse problems that does not require knowledge of the forward model. We develop efficient methods to compute an optimal low-rank regularized inverse matrix for a variety of error measures. In addition, we provide a natural framework for evaluating solutions and predicting uncertainties.

The paper is structured as follows. In Section 2 we provide background on regularization of inverse problems and discuss the connection to Bayes risk minimization. Some theoretical results are provided for the corresponding Bayes risk minimization problem. In Section 3 we describe numerical methods to efficiently solve (3), including a rank-update approach that can be used for various error measures and large-scale problems. Numerical results are presented in Section 4 and conclusions are provided in Section 5.

## 2 Background and Motivation

The singular value decomposition (SVD) is an important tool for analyzing regularization methods. For any real  $m \times n$  matrix  $\mathbf{A}$ , let  $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$  be the SVD, where  $\mathbf{\Sigma}$  is a diagonal matrix containing the singular values,  $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_{\text{rank}(\mathbf{A})} > 0$ , and orthogonal matrices  $\mathbf{U}$  and  $\mathbf{V}$  contain the left and right singular vectors  $\mathbf{u}_i, i = 1, 2, \dots, m$ , and  $\mathbf{v}_i, i = 1, 2, \dots, n$  respectively [16]. For a given positive integer  $r \leq \text{rank}(\mathbf{A})$ , define

$$\mathbf{A}_r = \mathbf{U}_r \mathbf{\Sigma}_r \mathbf{V}_r^\top,$$

where  $\mathbf{U}_r = [\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_r] \in \mathbb{R}^{m \times r}$ ,  $\mathbf{V}_r = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_r] \in \mathbb{R}^{n \times r}$ , and

$$\mathbf{\Sigma}_r = \begin{bmatrix} \sigma_1 & & & \\ & \sigma_2 & & \\ & & \ddots & \\ & & & \sigma_r \end{bmatrix} \in \mathbb{R}^{r \times r}.$$

Then, we denote  $\mathbf{A}_r^\dagger = \mathbf{V}_r \mathbf{\Sigma}_r^{-1} \mathbf{U}_r^\top$  to be the pseudoinverse of  $\mathbf{A}_r$  [17].

For ill-posed inverse problems, the singular values of  $\mathbf{A}$  decay to and cluster at zero. Spectral filtering methods provide solutions where components of the solution corresponding to the small singular values are damped or neglected. For example, the rank- $r$  truncated SVD (TSVD) solution can be written as

$$\boldsymbol{\xi}_{\text{TSVD}} = \mathbf{A}_r^\dagger \mathbf{b}. \quad (4)$$

In Figure 1, we provide TSVD sample mean errors for different values of  $r$  for a image deblurring problem. For TSVD, the choice of rank is crucial. For small rank, there are not enough singular vectors to capture the details of the solution, but for large rank, small singular values are included, and noise corrupts the solution. This is an illustrative example, so details are not specified here. See Section 4 for more thorough numerical investigations.

It is worth noting that the TSVD reconstruction matrix  $\mathbf{A}_r^\dagger$  is *one* example of a rank- $r$  matrix that solves (1). In this paper, we seek an *optimal* rank- $r$  matrix for all training data. Computing a solution to problem (3) is non-trivial, and in the next section we provide some underlying theory from a Bayes risk perspective.

### 2.1 Bayes risk minimization

By interpreting problem (3) as an empirical Bayes risk minimization problem [18, 13, 19], we can provide a statistical framework for analyzing the problem and motivate numerical algorithms for computing an optimal low-rank regularized inverse matrix. In the following we assume that  $\boldsymbol{\xi}$  and  $\boldsymbol{\delta}$  are random variables with given probability distributions  $\mathcal{P}_\xi$  and  $\mathcal{P}_\delta$ . Using the fact that  $\mathbf{Zb} - \boldsymbol{\xi} = (\mathbf{ZA} - \mathbf{I})\boldsymbol{\xi} + \mathbf{Z}\boldsymbol{\delta}$ , the *Bayes risk minimization problem* can be written as,

$$\tilde{\mathbf{Z}} = \arg \min_{\text{rank}(\mathbf{Z}) \leq r} \mathbb{E}_{\boldsymbol{\xi}, \boldsymbol{\delta}} (\rho((\mathbf{ZA} - \mathbf{I})\boldsymbol{\xi} + \mathbf{Z}\boldsymbol{\delta})), \quad (5)$$

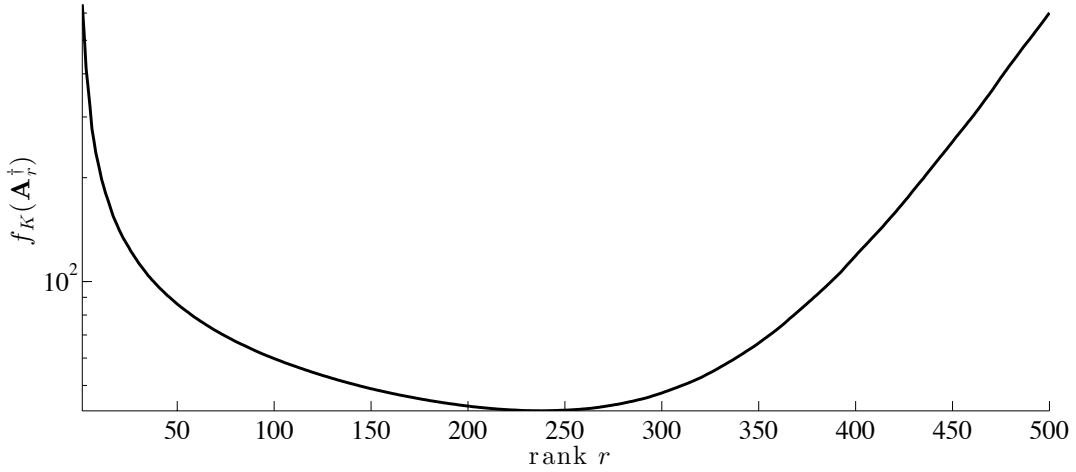


Figure 1: Illustration of a typical sample mean error  $f_K(\mathbf{A}_r^\dagger)$  on a logarithmic scale for Truncated SVD reconstructions for various ranks  $r$ .

where  $\mathbb{E}_{\boldsymbol{\xi}, \boldsymbol{\delta}}$  is the expected value with respect to the joint distribution of  $\boldsymbol{\xi}$  and  $\boldsymbol{\delta}$  [18, 13]. The particular choices of  $\rho$ ,  $\mathcal{P}_{\boldsymbol{\xi}}$ , and  $\mathcal{P}_{\boldsymbol{\delta}}$  in (5) determine an optimal reconstruction matrix  $\tilde{\mathbf{Z}}$ . Optimization problem (3) can be interpreted as an empirical Bayes risk minimization problem corresponding to Bayes risk minimization problem (5), where  $\boldsymbol{\xi}^{(k)}$  and  $\mathbf{b}^{(k)} = \mathbf{A}\boldsymbol{\xi}^{(k)} + \boldsymbol{\delta}^{(k)}$  denote realizations of the random variables  $\boldsymbol{\xi}$  and  $\boldsymbol{\delta}$ .

## 2.2 Bayes Risk Minimization for the 2-norm case

For the special case where  $\rho(\mathbf{x}) = \|\mathbf{x}\|_2^2$ , the *Bayes risk* can be written as

$$\tilde{f}(\mathbf{Z}) = \mathbb{E}_{\boldsymbol{\xi}, \boldsymbol{\delta}} (\|(\mathbf{Z}\mathbf{A} - \mathbf{I}_n)\boldsymbol{\xi} + \mathbf{Z}\boldsymbol{\delta}\|_2^2). \quad (6)$$

In this section, we provide closed form solutions for both the unconstrained and rank-constrained Bayes risk minimization problems for the 2-norm case.

We assume that the random variables  $\boldsymbol{\xi}$  and  $\boldsymbol{\delta}$  are statistically independent. Let  $\boldsymbol{\mu}_{\boldsymbol{\xi}}$  be the mean and  $\mathbf{C}_{\boldsymbol{\xi}}$  be the covariance matrix for the distribution  $\mathcal{P}_{\boldsymbol{\xi}}$ . Similarly,  $\boldsymbol{\mu}_{\boldsymbol{\delta}}$  and  $\mathbf{C}_{\boldsymbol{\delta}}$  are the mean and covariance matrix for  $\mathcal{P}_{\boldsymbol{\delta}}$ . We assume that both matrices are positive definite and that all entries are finite. One fact that will be useful to us is that the expectation of the quadratic form  $\boldsymbol{\phi}^\top \mathbf{K} \boldsymbol{\phi}$  involving a symmetric matrix  $\mathbf{K}$  and a random variable  $\boldsymbol{\phi}$  with covariance matrix  $\mathbf{C}_{\boldsymbol{\phi}}$  can be written as

$$\mathbb{E}_{\boldsymbol{\phi}}(\boldsymbol{\phi}^\top \mathbf{K} \boldsymbol{\phi}) = \mathbb{E}_{\boldsymbol{\phi}}(\boldsymbol{\phi})^\top \mathbf{K} \mathbb{E}_{\boldsymbol{\phi}}(\boldsymbol{\phi}) + \text{tr}(\mathbf{K} \mathbf{C}_{\boldsymbol{\phi}}),$$

where  $\text{tr}(\cdot)$  denotes the trace of a matrix [20]. With this framework and the Cholesky

decompositions  $\mathbf{C}_\xi = \mathbf{M}_\xi \mathbf{M}_\xi^\top$  and  $\mathbf{C}_\delta = \mathbf{M}_\delta \mathbf{M}_\delta^\top$ , we can write the Bayes risk (6) as,

$$\begin{aligned}\tilde{f}(\mathbf{Z}) &= \mathbb{E}_{\xi, \delta} (\|(\mathbf{Z}\mathbf{A} - \mathbf{I}_n)\xi + \mathbf{Z}\delta\|_2^2) \\ &= \mathbb{E}_\xi (\xi^\top (\mathbf{Z}\mathbf{A} - \mathbf{I}_n)^\top (\mathbf{Z}\mathbf{A} - \mathbf{I}_n) \xi) + 2 \mathbb{E}_\delta (\delta^\top) \mathbf{Z}^\top (\mathbf{Z}\mathbf{A} - \mathbf{I}_n) \mathbb{E}_\xi (\xi) + \mathbb{E}_\delta (\delta^\top \mathbf{Z}^\top \mathbf{Z} \delta) \\ &= \mu_\xi^\top (\mathbf{Z}\mathbf{A} - \mathbf{I}_n)^\top (\mathbf{Z}\mathbf{A} - \mathbf{I}_n) \mu_\xi + \text{tr}((\mathbf{Z}\mathbf{A} - \mathbf{I}_n)^\top (\mathbf{Z}\mathbf{A} - \mathbf{I}_n) \mathbf{M}_\xi \mathbf{M}_\xi^\top) \\ &\quad + 2 \mu_\delta^\top \mathbf{Z}^\top (\mathbf{Z}\mathbf{A} - \mathbf{I}_n) \mu_\xi + \mu_\delta^\top \mathbf{Z}^\top \mathbf{Z} \mu_\delta + \text{tr}(\mathbf{Z}^\top \mathbf{Z} \mathbf{M}_\delta \mathbf{M}_\delta^\top) \\ &= \|(\mathbf{Z}\mathbf{A} - \mathbf{I}_n) \mu_\xi\|_2^2 + \|(\mathbf{Z}\mathbf{A} - \mathbf{I}_n) \mathbf{M}_\xi\|_F^2 + 2 \mu_\delta^\top \mathbf{Z}^\top (\mathbf{Z}\mathbf{A} - \mathbf{I}_n) \mu_\xi + \|\mathbf{Z} \mu_\delta\|_2^2 + \|\mathbf{Z} \mathbf{M}_\delta\|_F^2,\end{aligned}$$

where  $\|\cdot\|_F$  denotes the Frobenius norm of a matrix.

We can simplify this final expression by *pre-whitening* the noise  $\delta$ , thereby converting to a coordinate system where the noise mean is zero and the noise covariance is a multiple of the identity. Define the random variable

$$\chi = \mathbf{M}_\delta^{-1}(\delta - \mu_\delta). \quad (7)$$

Then,  $\mu_\chi = \mathbf{0}$  and  $\mathbf{C}_\chi = \mathbf{I}_m$ , and we have

$$\mathbf{M}_\delta^{-1} \mathbf{A} \xi + \chi = \mathbf{M}_\delta^{-1}(\mathbf{b} - \mu_\delta), \quad (8)$$

which can be written as,

$$\tilde{\mathbf{A}} \xi + \chi = \tilde{\mathbf{b}}, \quad (9)$$

where  $\tilde{\mathbf{A}} = \mathbf{M}_\delta^{-1} \mathbf{A}$  and  $\tilde{\mathbf{b}} = \mathbf{M}_\delta^{-1}(\mathbf{b} - \mu_\delta)$ . Thus, without loss of generality (although with the need for numerical care if  $\mathbf{M}_\delta$  is ill-conditioned), we can always assume that  $\mu_\delta = \mathbf{0}$  and  $\mathbf{C}_\delta = \eta^2 \mathbf{I}_m$ . Therefore, our Bayes risk (6) takes the form

$$\tilde{f}(\mathbf{Z}) = \|(\mathbf{Z}\mathbf{A} - \mathbf{I}_n) \mu_\xi\|_2^2 + \|(\mathbf{Z}\mathbf{A} - \mathbf{I}_n) \mathbf{M}_\xi\|_F^2 + \eta^2 \|\mathbf{Z}\|_F^2. \quad (10)$$

Under certain conditions, summarized in the following theorem, there exists a unique global minimizer for (10).

**Theorem 2.1.** *Given  $\mathbf{A} \in \mathbb{R}^{m \times n}$ ,  $\mathbf{M}_\xi \in \mathbb{R}^{n \times n}$ ,  $\mu_\xi \in \mathbb{R}^n$  and  $\eta \in \mathbb{R}$ , consider the problem of finding a matrix  $\check{\mathbf{Z}} \in \mathbb{R}^{n \times m}$  with*

$$\check{\mathbf{Z}} = \arg \min_{\mathbf{Z}} \left( \|(\mathbf{Z}\mathbf{A} - \mathbf{I}_n) \mu_\xi\|_2^2 + \|(\mathbf{Z}\mathbf{A} - \mathbf{I}_n) \mathbf{M}_\xi\|_F^2 + \eta^2 \|\mathbf{Z}\|_F^2 \right). \quad (11)$$

*If either  $\eta \neq 0$  or if  $\mathbf{A}$  has rank  $m$  and  $\mu_\xi \mu_\xi^\top + \mathbf{M}_\xi \mathbf{M}_\xi^\top$  is nonsingular, then the unique global minimizer is*

$$\check{\mathbf{Z}} = (\mu_\xi \mu_\xi^\top + \mathbf{M}_\xi \mathbf{M}_\xi^\top) \mathbf{A}^\top [\mathbf{A} (\mu_\xi \mu_\xi^\top + \mathbf{M}_\xi \mathbf{M}_\xi^\top) \mathbf{A}^\top + \eta^2 \mathbf{I}_m]^{-1}. \quad (12)$$

*Proof.* Using derivatives of norms and traces, we get

$$\frac{d\tilde{f}(\mathbf{Z})}{d\mathbf{Z}} = 2(\mathbf{Z}\mathbf{A} - \mathbf{I}_n)\boldsymbol{\mu}_\xi\boldsymbol{\mu}_\xi^\top\mathbf{A}^\top + 2(\mathbf{Z}\mathbf{A} - \mathbf{I}_n)\mathbf{M}_\xi\mathbf{M}_\xi^\top\mathbf{A}^\top + 2\eta^2\mathbf{Z}. \quad (13)$$

Setting the above derivative equal to zero gives the following first order condition,

$$\mathbf{Z}[\mathbf{A}(\boldsymbol{\mu}_\xi\boldsymbol{\mu}_\xi^\top + \mathbf{M}_\xi\mathbf{M}_\xi^\top)\mathbf{A}^\top + \eta^2\mathbf{I}_m] = (\boldsymbol{\mu}_\xi\boldsymbol{\mu}_\xi^\top + \mathbf{M}_\xi\mathbf{M}_\xi^\top)\mathbf{A}^\top. \quad (14)$$

Under our assumptions, the matrix  $\mathbf{A}(\boldsymbol{\mu}_\xi\boldsymbol{\mu}_\xi^\top + \mathbf{M}_\xi\mathbf{M}_\xi^\top)\mathbf{A}^\top + \eta^2\mathbf{I}_m$  is invertible, so we get the unique solution  $\check{\mathbf{Z}}$  that is given in the statement of the Theorem. It is a minimizer since the second derivative matrix of  $\tilde{f}(\mathbf{Z})$  is positive definite.  $\square$

To summarize, we have minimized the Bayes risk (6) under the assumptions that  $\boldsymbol{\xi}$  and  $\boldsymbol{\delta}$  are statistically independent random variables with symmetric positive definite finite covariance matrices.

Next we consider the rank-constrained Bayes risk minimization problem. With the additional assumption that  $\boldsymbol{\mu}_\xi = \mathbf{0}$ , problem (5) can be written as

$$\min_{\text{rank}(\mathbf{Z}) \leq r} \|(\mathbf{Z}\mathbf{A} - \mathbf{I})\mathbf{M}\|_F^2 + \alpha^2 \|\mathbf{Z}\|_F^2, \quad (15)$$

where  $\mathbf{M} = \mathbf{M}_\xi$  and  $\alpha = \eta$ . A closed form solution to (15) was recently provided in [21], and the result is summarized in the following theorem.

**Theorem 2.2.** *Given a matrix  $\mathbf{A} \in \mathbb{R}^{m \times n}$  where  $\text{rank}(\mathbf{A}) \leq n \leq m$  and an invertible matrix  $\mathbf{M} \in \mathbb{R}^{n \times n}$ , let their generalized singular value decomposition be  $\mathbf{A} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{G}^{-1}$ ,  $\mathbf{M} = \mathbf{G}\mathbf{S}^{-1}\mathbf{V}^\top$ . Let  $\alpha$  be a given parameter, nonzero if  $\text{rank}(\mathbf{A}) < m$ . Let  $r \leq \text{rank}(\mathbf{A})$  be a given positive integer. Define  $\mathbf{D} = \boldsymbol{\Sigma}\mathbf{S}^{-2}\boldsymbol{\Sigma}^\top + \alpha^2\mathbf{I}$ . Let the symmetric matrix  $\mathbf{H} = \mathbf{G}\mathbf{S}^{-4}\boldsymbol{\Sigma}^\top\mathbf{D}^{-1}\boldsymbol{\Sigma}\mathbf{G}^\top$  have eigenvalue decomposition  $\mathbf{H} = \check{\mathbf{V}}\boldsymbol{\Lambda}\check{\mathbf{V}}^\top$  with eigenvalues ordered so that  $\lambda_j \geq \lambda_i$  for  $j < i \leq n$ . Then a global minimizer  $\check{\mathbf{Z}} \in \mathbb{R}^{n \times m}$  of problem (15) is*

$$\check{\mathbf{Z}} = \check{\mathbf{V}}_r\check{\mathbf{V}}_r^\top\mathbf{G}\mathbf{S}^{-2}\boldsymbol{\Sigma}^\top\mathbf{D}^{-1}\mathbf{U}^\top, \quad (16)$$

where  $\check{\mathbf{V}}_r$  contains the first  $r$  columns of  $\check{\mathbf{V}}$ . Moreover this  $\check{\mathbf{Z}}$  is the unique global minimizer of (15) if and only if  $\lambda_r > \lambda_{r+1}$ .

Next we consider the case where  $\mathbf{C}_\xi = \beta^2\mathbf{I}_n$  for some scalar  $\beta$ . In this case the Bayes risk simplifies to

$$\tilde{f}(\mathbf{Z}) = \beta^2 \|\mathbf{Z}\mathbf{A} - \mathbf{I}_n\|_F^2 + \eta^2 \|\mathbf{Z}\|_F^2, \quad (17)$$

and, by rescaling, the low-rank Bayes risk minimization problem can be restated as optimization problem (15) where  $\alpha^2 = \frac{\eta^2}{\beta^2}$  represents the noise-to-signal ratio and  $\mathbf{M} = \mathbf{I}_n$ . It was shown in [21] that for this problem, the global minimizer simplifies to

$$\check{\mathbf{Z}} = \mathbf{V}_r\boldsymbol{\Psi}_r\mathbf{U}_r^\top, \quad (18)$$

where  $\mathbf{\Psi}_r = \text{diag}\left(\frac{\sigma_1}{\sigma_1^2 + \alpha^2}, \dots, \frac{\sigma_r}{\sigma_r^2 + \alpha^2}\right)$  and  $\mathbf{V}_r$  and  $\mathbf{U}_r$  were defined in the beginning of this section. Moreover,  $\tilde{\mathbf{Z}}$  is the *unique* global minimizer if and only if  $\sigma_r > \sigma_{r+1}$ . For problems where  $\mathbf{A}$  is known, matrix  $\tilde{\mathbf{Z}}$  can be used to compute a filtered solution to inverse problems. That is,

$$\mathbf{x}_{\text{filter}} \equiv \tilde{\mathbf{Z}}\mathbf{b} = \sum_{j=1}^n \phi_j \frac{\mathbf{u}_j^\top \mathbf{b}}{\sigma_j} \mathbf{v}_j,$$

where the filter factors are given by

$$\phi_j = \begin{cases} \frac{\sigma_j^2}{\sigma_j^2 + \alpha^2}, & \text{for } j \leq r, \\ 0, & \text{for } j > r. \end{cases}$$

Here,  $\tilde{\mathbf{Z}}$  is the *truncated Tikhonov* reconstruction matrix [21].

### 3 Computing Low-rank Regularized Inverse Matrices

In this section, we describe numerical methods to efficiently compute a low-rank regularized inverse matrix, i.e., a solution to (3), for problems where matrix  $\mathbf{A}$  is unknown and training data are available. We first describe a method that can be used for the special case where  $\rho = \|\cdot\|_2^2$ . Further we discuss a rank-update approach that can be utilized for large-scale problems and for general error measures.

#### 3.1 2-norm case

For the case  $\rho = \|\cdot\|_2^2$ , we can use the relationship between the 2-norm of a vector and the Frobenius norm of a matrix to rewrite optimization problem (3) as

$$\hat{\mathbf{Z}} = \arg \min_{\text{rank}(\mathbf{Z}) \leq r} \frac{1}{K} \sum_{k=1}^K \left\| \mathbf{Z}\mathbf{b}^{(k)} - \boldsymbol{\xi}^{(k)} \right\|_2^2 = \frac{1}{K} \left\| \mathbf{Z}\mathbf{B} - \mathbf{C} \right\|_{\text{F}}^2, \quad (19)$$

where  $\mathbf{B} = [\mathbf{b}^{(1)}, \dots, \mathbf{b}^{(K)}]$  and  $\mathbf{C} = [\boldsymbol{\xi}^{(1)}, \dots, \boldsymbol{\xi}^{(K)}]$ . A closed form solution for (19) exists, and the result is summarized in the following theorem.

**Theorem 3.1.** *Let matrices  $\mathbf{B} \in \mathbb{R}^{m \times K}$  and  $\mathbf{C} \in \mathbb{R}^{n \times K}$  be given. Further let  $\tilde{\mathbf{V}}$  be the matrix of right singular vectors of  $\mathbf{B}$  and define  $\mathbf{P} = \mathbf{C}\tilde{\mathbf{V}}_s(\tilde{\mathbf{V}}_s)^\top$  where  $s = \text{rank}(\mathbf{B})$ . Then*

$$\hat{\mathbf{Z}} = \mathbf{P}_r \mathbf{B}^\dagger$$

*is a solution to the minimization problem (19). This solution is unique if and only if either  $r \geq \text{rank}(\mathbf{P})$  or  $1 \leq r \leq \text{rank}(\mathbf{P})$  and  $\bar{\sigma}_r > \bar{\sigma}_{r+1}$ , where  $\bar{\sigma}_r$  and  $\bar{\sigma}_{r+1}$  denote the  $r$  and  $(r+1)$ st singular values of  $\mathbf{P}$ .*

*Proof.* This result is a special case of the main result of Sondermann [22] and was reproved with if-and-only-if uniqueness conditions in Theorem 2.1 of Friedland and Torokhti [23].  $\square$

The main computational cost of constructing  $\hat{\mathbf{Z}}$  using Theorem 3.1 lies in two SVD computations (for matrices  $\mathbf{B}$  and  $\mathbf{P}$ ). If these computations are possible and minimizing the mean squared error is desired, then we propose to use the above result to compute optimal low-rank regularized inverse matrix  $\hat{\mathbf{Z}}$ , which in turn can be used to solve inverse problems [24].

### 3.2 A rank-update approach

For general error measures or for large problems with large training sets, we propose a rank-update approach for obtaining optimal low-rank regularized inverse matrix  $\hat{\mathbf{Z}}$ . The approach is motivated by the following corollary to Theorem 2.2.

**Corollary.** *Assume all conditions of Theorem 2.2 are fulfilled. Let  $\check{\mathbf{Z}}^J$  be a global minimizer of (15) for maximal rank  $J$  and  $\check{\mathbf{Z}}^{J+\ell}$  be a global minimizer of (15) for maximal rank  $J + \ell$ . Then  $\bar{\mathbf{Z}}^\ell = \check{\mathbf{Z}}^{J+\ell} - \check{\mathbf{Z}}^J$  is the global minimizer of*

$$\min_{\text{rank}(\mathbf{Z}) \leq \ell} \|((\check{\mathbf{Z}}^J + \mathbf{Z})\mathbf{A} - \mathbf{I})\mathbf{M}\|_{\text{F}}^2 + \alpha^2 \|\check{\mathbf{Z}}^J + \mathbf{Z}\|_{\text{F}}^2. \quad (20)$$

Furthermore,  $\bar{\mathbf{Z}}^\ell$  is the unique global minimizer if and only if  $\lambda_J > \lambda_{J+1}$  and  $\lambda_{J+\ell} > \lambda_{J+\ell+1}$  with  $J + \ell + 1 < \text{rank}(\mathbf{A})$ .

Corollary 3.2 states that in the Bayes framework for the 2-norm case, a rank-update approach (in exact arithmetic) can be used to compute an optimal low-rank regularized matrix. That is, assume we are given the rank- $J$  approximation  $\check{\mathbf{Z}}^J$  and we are interested in finding  $\check{\mathbf{Z}}^{J+\ell}$ . To calculate the optimal rank- $(J+\ell)$  approximation  $\check{\mathbf{Z}}^{J+\ell}$ , we just need to solve a rank- $\ell$  optimization problem of the form (20) and then update the solution,  $\check{\mathbf{Z}}^{J+\ell} = \check{\mathbf{Z}}^J + \bar{\mathbf{Z}}^\ell$ .

We are interested in the empirical Bayes problem, and we extend the rank-update approach to problem (3), where the basic idea is to iteratively build up matrices  $\hat{\mathbf{X}} \in \mathbb{R}^{n \times r}$  and  $\hat{\mathbf{Y}} \in \mathbb{R}^{m \times r}$  such that  $\hat{\mathbf{Z}} = \hat{\mathbf{X}}\hat{\mathbf{Y}}^\top$ . Suppose we are given the optimal rank- $J$  matrix  $\hat{\mathbf{Z}}$ . That is,  $\hat{\mathbf{Z}}$  solves (3) for  $r = J$ . Then, we compute a rank- $\ell$  update to  $\hat{\mathbf{Z}}$  by solving the following rank- $\ell$  optimization problem,

$$(\hat{\mathbf{X}}^\ell, \hat{\mathbf{Y}}^\ell) = \arg \min_{\mathbf{X} \in \mathbb{R}^{n \times \ell}, \mathbf{Y} \in \mathbb{R}^{m \times \ell}} f_K(\mathbf{X}, \mathbf{Y}) = \frac{1}{K} \sum_{k=1}^K \rho((\hat{\mathbf{Z}} + \mathbf{X}\mathbf{Y}^\top)\mathbf{b}^{(k)} - \boldsymbol{\xi}^{(k)}), \quad (21)$$

and a rank- $(J + \ell)$  regularized inverse matrix can be computed as  $\hat{\mathbf{Z}} = \hat{\mathbf{Z}} + \hat{\mathbf{X}}^\ell(\hat{\mathbf{Y}}^\ell)^\top$ . This process continues until the desired rank is achieved. A summary of the general algorithm can be found below, see Algorithm 1.

Next we provide some computational remarks for Algorithm 1.



---

**Algorithm 1** Rank- $\ell$  update approach for computing a solution to (3)

---

- 1: Initialize  $\widehat{\mathbf{Z}} = \mathbf{0}_{n \times m}$ ,  $\widehat{\mathbf{X}} = [\ ]$ ,  $\widehat{\mathbf{Y}} = [\ ]$
  - 2: **while** stopping criteria not satisfied **do**
  - 3:   solve rank- $\ell$  update problem (21)
  - 4:   update matrix inverse approximation:  $\widehat{\mathbf{Z}} \leftarrow \widehat{\mathbf{Z}} + \widehat{\mathbf{X}}^\ell (\widehat{\mathbf{Y}}^\ell)^\top$
  - 5:   update solutions:  $\widehat{\mathbf{X}} \leftarrow [\widehat{\mathbf{X}}, \widehat{\mathbf{X}}^\ell]$ ,  $\widehat{\mathbf{Y}} \leftarrow [\widehat{\mathbf{Y}}, \widehat{\mathbf{Y}}^\ell]$
  - 6: **end while**
- 

**Choice of  $\ell$ .**

At each iteration of Algorithm 1, we must solve optimization problem (21), which has  $(m + n)\ell$  unknowns. Hence, the choice of  $\ell$  will depend on the size of the problem and may even vary per iteration, e.g.,  $\ell_1, \ell_2, \dots$ . For large-scale problems such as those arising in image reconstruction, we recommend updates of size  $\ell = 1$  to keep optimization problem (21) as small as possible.

**Stopping criteria.**

Stopping criteria for Algorithm 1 should have four components. First, a safe-guard maximal rank condition. That is, we stop if  $\text{rank}(\widehat{\mathbf{Z}})$  is greater than some pre-determined integer value  $r_{\max}$ . Secondly, let  $f_{\max}$  be a user-defined error tolerance for the function value, then we stop if  $f_K \leq f_{\max}$ . Similarly, we can monitor the relative improvement of  $f_K$  in each iteration and stop if the relative improvement is smaller than some tolerance. Lastly, if a rank- $\ell$  update does not substantially improve  $f_K$ , the iteration should be terminated. More specifically, let  $\widehat{\mathbf{X}}^\ell = [\widehat{\mathbf{x}}_1, \dots, \widehat{\mathbf{x}}_\ell]$  and choose tolerance  $\tau$ . Then we can control the rank deficiency of  $\mathbf{X}^\ell$  by stopping if  $\min\{\|\widehat{\mathbf{x}}_1\|_\infty, \dots, \|\widehat{\mathbf{x}}_\ell\|_\infty\} \leq \tau$ . The particular choices of stopping criteria and tolerances are problem dependent.

**Solving rank- $\ell$  update problem (21).**

The main challenge for Algorithm 1 is in line 3, where we need to solve the rank- $\ell$  update problem (21). First, we remark that the solution  $(\widehat{\mathbf{X}}^\ell, \widehat{\mathbf{Y}}^\ell)$  is not unique. In fact,  $\widehat{\mathbf{X}}^\ell (\widehat{\mathbf{Y}}^\ell)^\top = \frac{1}{a} \widehat{\mathbf{X}}^\ell (a \widehat{\mathbf{Y}}^\ell)^\top$  for any  $a \neq 0$ . As a consequence optimization problem (21) is ill-posed and a particular solution needs to be selected. Various computational approaches exist to handle this kind of ill-posedness. One way to tackle this problem is to constrain each column in  $\mathbf{X}$  to have norm 1, i.e.,  $\|\mathbf{X}_{(:,j)}\|_2 = 1$  for  $j = 1, \dots, \ell$ , and solve the constrained or relaxed joint optimization problem. We found this approach to be computationally inefficient and instead propose an alternating optimization approach. Notice that for any convex function  $\rho$ , function  $f_K(\mathbf{X}, \mathbf{Y})$  is convex in  $\mathbf{X}$  and convex in  $\mathbf{Y}$ , so an alternating optimization approach avoids the above ill-posedness. Until convergence, the solution to  $\min_{\mathbf{Y}} f_K(\mathbf{X}, \mathbf{Y})$  is computed for fixed  $\mathbf{X}$ , and then the solution to  $\min_{\mathbf{X}} f_K(\mathbf{X}, \mathbf{Y})$  is computed for fixed  $\mathbf{Y}$ . Columns of  $\mathbf{X}$  are normalized at each iteration. The alternating optimization is summarized

in Algorithm 2.

---

**Algorithm 2** Alternating optimization for solving (21)

---

- 1: choose  $\mathbf{X}^{(0)} \in \mathbb{R}^{m \times \ell}$ ,  $i = 0$
  - 2: **while** stopping criteria not satisfied **do**
  - 3:   Solve  $\mathbf{Y}^{(i+1)} = \arg \min_{\mathbf{Y}} f_K(\mathbf{X}^{(i)}, \mathbf{Y})$
  - 4:   Solve  $\mathbf{X}^{(i+1)} = \arg \min_{\mathbf{X}} f_K(\mathbf{X}, \mathbf{Y}^{(i+1)})$
  - 5:    $\mathbf{X}_{(:,j)}^{(i+1)} \leftarrow \mathbf{X}_{(:,j)}^{(i+1)} / \left\| \mathbf{X}_{(:,j)}^{(i+1)} \right\|_2$  for  $j = 1, \dots, \ell$
  - 6:    $i \leftarrow i + 1$
  - 7: **end while**
- 

Methods for solving optimization problems in lines 3 and 4 of Algorithm 2 depend on the specific form of  $\rho$ . For instance for  $\rho = \|\cdot\|_p^p$  with  $2 \leq p < \infty$  we suggest using Gauss-Newton type methods. For  $p$ -norms with  $1 < p < 2$  we propose to use Quasi-Newton type methods such as SR1, BFGS, and LBFGS [25] or utilize smoothing techniques to compensate for discontinuities in the derivative. Standard stopping criteria can be used for these optimization problems, as well as for Algorithm 2 [26]. The required gradients of  $f_K$  and Gauss-Newton approximation to the Hessian are provided below.

Let  $\mathbf{e}^{(k)} = \mathbf{X}\mathbf{Y}^\top \mathbf{b}^{(k)} - \mathbf{c}^{(k)}$ ,  $\mathbf{x} = \text{vec}(\mathbf{X})$ , and  $\mathbf{y} = \text{vec}(\mathbf{Y})$ . Then the gradients are given by

$$\nabla_{\mathbf{x}} f_K = \frac{1}{K} (\mathbf{J}_{\mathbf{x}})^\top \boldsymbol{\rho} \quad \text{and} \quad \nabla_{\mathbf{y}} f_K = \frac{1}{K} (\mathbf{J}_{\mathbf{y}})^\top \boldsymbol{\rho}, \quad (22)$$

where  $\boldsymbol{\rho} = \text{vec}([\boldsymbol{\rho}^{(1)}, \dots, \boldsymbol{\rho}^{(K)}])$  with  $\boldsymbol{\rho}^{(k)} = \nabla \rho$  evaluated at  $\mathbf{e}^{(k)}$ . The Jacobian matrices have the form,

$$\mathbf{J}_{\mathbf{x}} = \begin{bmatrix} \mathbf{J}_{\mathbf{x}_1}^{(1)} & \dots & \mathbf{J}_{\mathbf{x}_\ell}^{(1)} \\ \vdots & & \vdots \\ \mathbf{J}_{\mathbf{x}_1}^{(K)} & \dots & \mathbf{J}_{\mathbf{x}_\ell}^{(K)} \end{bmatrix} \quad \text{and} \quad \mathbf{J}_{\mathbf{y}} = \begin{bmatrix} \mathbf{J}_{\mathbf{y}_1}^{(1)} & \dots & \mathbf{J}_{\mathbf{y}_\ell}^{(1)} \\ \vdots & & \vdots \\ \mathbf{J}_{\mathbf{y}_1}^{(K)} & \dots & \mathbf{J}_{\mathbf{y}_\ell}^{(K)} \end{bmatrix},$$

where  $\mathbf{J}_{\mathbf{x}_j}^{(k)} = \mathbf{y}_j^\top \mathbf{b}^{(k)} \mathbf{I}_n$  and  $\mathbf{J}_{\mathbf{y}_j}^{(k)} = \mathbf{x}_j (\mathbf{b}^{(k)})^\top$ . The Gauss-Newton approximations of the Hessians are given by

$$\mathbf{H}_{\mathbf{x}} = \frac{1}{K} (\mathbf{J}_{\mathbf{x}})^\top \mathbf{D} \mathbf{J}_{\mathbf{x}} \quad \text{and} \quad \mathbf{H}_{\mathbf{y}} = \frac{1}{K} (\mathbf{J}_{\mathbf{y}})^\top \mathbf{D} \mathbf{J}_{\mathbf{y}}, \quad (23)$$

where  $\mathbf{D} = \text{diag}(\mathbf{D}^{(1)}, \dots, \mathbf{D}^{(K)})$  and  $\mathbf{D}^{(k)} = \nabla^2 \rho$  evaluated at  $\mathbf{e}^{(k)}$ . Note that  $\mathbf{D}$  is diagonal for any error measures  $\rho$  of the form  $\rho(\mathbf{x}) = \sum_{i=1}^n \phi(x_i)$  which includes the Huber function and all  $p$ -norms with  $p < \infty$ .

Another important remark is that for computational efficiency, we do not need to form  $\hat{\mathbf{Z}}$  explicitly. That is, if we let  $\mathbf{c}^{(k)} = \boldsymbol{\xi}^{(k)} - \hat{\mathbf{Z}}\mathbf{b}^{(k)}$ , then (21) can be reformulated as

$$(\hat{\mathbf{X}}^\ell, \hat{\mathbf{Y}}^\ell) = \arg \min_{\mathbf{X} \in \mathbb{R}^{n \times \ell}, \mathbf{Y}^{m \times \ell}} \frac{1}{K} \sum_{k=1}^K \rho(\mathbf{X}\mathbf{Y}^\top \mathbf{b}^{(k)} - \mathbf{c}^{(k)}), \quad (24)$$

where vectors  $\mathbf{c}^{(k)}$  should be updated at each iteration as  $\mathbf{c}^{(k)} \leftarrow \mathbf{c}^{(k)} - \hat{\mathbf{X}}^\ell (\hat{\mathbf{Y}}^\ell)^\top \mathbf{b}^{(k)}$ .

### Using $\hat{\mathbf{Z}}$ to solve inverse problems.

We have described some numerical approaches for computing  $\hat{\mathbf{Z}}$ , and it is worthwhile to note that all of these computations can be done off-line. Once computed, the optimal low-rank regularized inverse matrix  $\hat{\mathbf{Z}}$  can be used to solve inverse problems very quickly, requiring only one matrix-vector multiplication  $\hat{\mathbf{Z}}\mathbf{b}$ , if  $\hat{\mathbf{Z}}$  is obtained explicitly, or two multiplications  $\hat{\mathbf{X}}(\hat{\mathbf{Y}}^\top \mathbf{b})$  if  $\hat{\mathbf{X}}$  and  $\hat{\mathbf{Y}}$  are computed separately. This is particularly important for applications where hundreds of inverse problems need to be solved in real time. The ability to obtain fast and accurate reconstructions of many inverse problems in real time quickly compensates for the initial cost of constructing  $\hat{\mathbf{Z}}$  (or  $\hat{\mathbf{X}}$  and  $\hat{\mathbf{Y}}$ ).

One of the significant advantages of our approach is that we seek a regularized inverse matrix  $\hat{\mathbf{Z}}$  directly, so that knowledge of the forward model  $\mathbf{A}$  is not required. However, for applications where  $\mathbf{A}$  is desired, our framework could be equivalently used to first estimate a low-rank matrix  $\mathbf{A}$  from the training data by solving Equation (3) where  $\mathbf{Z}$  is  $\mathbf{A}$ , and  $\mathbf{b}^k$  and  $\boldsymbol{\xi}^k$  are switched. This approach is similar to ideas from machine learning [27, 28] and are related to matrix completion problems where the goal is to reconstruct unknown or missing values in the forward matrix [29]. Then, once  $\hat{\mathbf{A}}$  is computed, standard regularization methods may be required to solve inverse problems. If this is the case, we propose to use the training data to obtain optimal regularization parameters, as discussed in [14]. Since the ultimate goal in many applications is to compute solutions to inverse problems, obtaining an approximation of  $\mathbf{A}$  is not necessary and may even introduce additional errors. Thus, if  $\mathbf{A}$  is not required, our approach to compute the optimal regularized inverse matrix is more efficient and less prone to errors.

## 4 Numerical Results

In this section, we compute optimal low-rank regularized inverse matrices for a variety of problems and evaluate the performance of these matrices when solving inverse problems. We generated a training set of  $K$  observations,  $\boldsymbol{\xi}^{(k)}$  and  $\mathbf{b}^{(k)}$ ,  $k = 1, \dots, K$ . Then for various error measures,  $\rho$ , we compute optimal low-rank regularized inverse matrices  $\hat{\mathbf{Z}}$  as described in Section 3. To validate the performance of the computed optimal low-rank regularized inverse matrices, we generate a validation set of  $\bar{K}$  different observations,  $\bar{\mathbf{b}}^{(k)}$ ,  $k = 1, \dots, \bar{K}$ , and compute the error between the reconstructed solution and the true solution  $\bar{\boldsymbol{\xi}}^{(k)}$ ,  $k = 1, \dots, \bar{K}$ , i.e.,  $e_k(\hat{\mathbf{Z}}) = \rho(\hat{\mathbf{Z}}\bar{\mathbf{b}}^{(k)} - \bar{\boldsymbol{\xi}}^{(k)})$ . For each experiment, we report the sample mean error  $f_K(\hat{\mathbf{Z}})$  as computed in (2), along with other descriptive statistics regarding the error distributions.

Standard regularization methods such as TSVD require knowledge of  $\mathbf{A}$ , while our approach requires training data. Thus, it is difficult to provide a fair comparison of methods. In our experiments, we compare our reconstructions to two “standard” reconstructions:

- **TSVD- $\mathbf{A}$ :** This refers to the TSVD reconstruction, as computed in (4), and it requires matrix  $\mathbf{A}$  and its SVD. The rank- $r$  reconstruction matrix for TSVD- $\mathbf{A}$  is given by  $\mathbf{Z} = \mathbf{A}_r^\dagger$ .
- **TSVD- $\hat{\mathbf{A}}$ :** If  $\mathbf{A}$  is not available, one could estimate it from the training data by solving the following rank-constrained optimization problem,

$$\hat{\mathbf{A}} = \arg \min_{\text{rank}(\mathbf{A}) \leq \bar{r}} \frac{1}{K} \sum_{k=1}^K \rho \left( \mathbf{A} \boldsymbol{\xi}^{(k)} - \mathbf{b}^{(k)} \right), \quad (25)$$

for some fixed integer  $\bar{r}$ . We propose to use methods described in Section 3 to obtain matrix  $\hat{\mathbf{A}}$ . Then, the computed matrix  $\hat{\mathbf{A}}$  could be used to construct a TSVD reconstruction. We refer to this approach as TSVD- $\hat{\mathbf{A}}$ , and the rank- $r$  reconstruction matrix for TSVD- $\hat{\mathbf{A}}$  for  $r \leq \bar{r}$  is given by  $\mathbf{Z} = \hat{\mathbf{A}}_r^\dagger$ . This seems to be a fair comparison to our approach since only training data is used; however, a potential disadvantage of TSVD- $\hat{\mathbf{A}}$  is that  $\hat{\mathbf{A}}$  can be large, and computing its SVD may be prohibitive.

We present results for three experiments: Experiments 1 and 2 are 1D examples, and Experiment 3 is a 2D example.

## 4.1 One-dimensional examples

For the 1D examples, we use random time signals by generating piecewise constant functions as seen in Figure 2. Here, the number of jumps in each signal is a randomly chosen integer between 3 and 20, drawn from a uniform distribution. The times for the jumps and the signal strength were both randomly selected in the interval  $[0, 1]$ , also from a uniform distribution. The discretized signal has length  $n = 150$ . The corresponding blurred signals were computed as in Equation (1), where matrix  $\mathbf{A}$  is  $150 \times 150$  and represents convolution with a 1D Gaussian kernel with zero mean and variance 2, and  $\boldsymbol{\delta}$  represents Gaussian white noise. The noise level for the blurred signals was 0.01, which means that  $\|\boldsymbol{\delta}\|_2^2 / \|\mathbf{A}\boldsymbol{\xi}\|_2^2 \approx 0.01$ . Training and validation signals were generated in the same manner. Figure 2 displays four sample signals from the validation set and their corresponding observed signals.

**Experiment 1.** For this experiment, we consider the case where  $\rho = \frac{1}{2} \|\cdot\|_2^2$  and compute an *optimal regularized inverse matrix* (ORIM) as described in Section 3.1. We refer to this approach as ORIM<sub>2</sub> and we refer to the computed matrix as  $\hat{\mathbf{Z}}_2$ .

First, we fix the number of training signals to  $K = 10,000$ . For varying ranks  $r$ , we compute the sample mean error for ORIM<sub>2</sub>, TSVD- $\mathbf{A}$ , and TSVD- $\hat{\mathbf{A}}$  for a set of 10,000 different validation signals. For computing  $\hat{\mathbf{A}}$ ,  $\bar{r}$  was fixed to 100. Errors plots as a function of rank are shown in Figure 3. We observe that the sample mean error for ORIM<sub>2</sub> is monotonically decreasing and consistently smaller than the sample mean errors for both TSVD- $\mathbf{A}$  and TSVD- $\hat{\mathbf{A}}$ . This is true for all ranks. Furthermore, both TSVD- $\mathbf{A}$  and TSVD- $\hat{\mathbf{A}}$  exhibit semi-convergence behavior similar to Figure 1, whereby the error initially improves

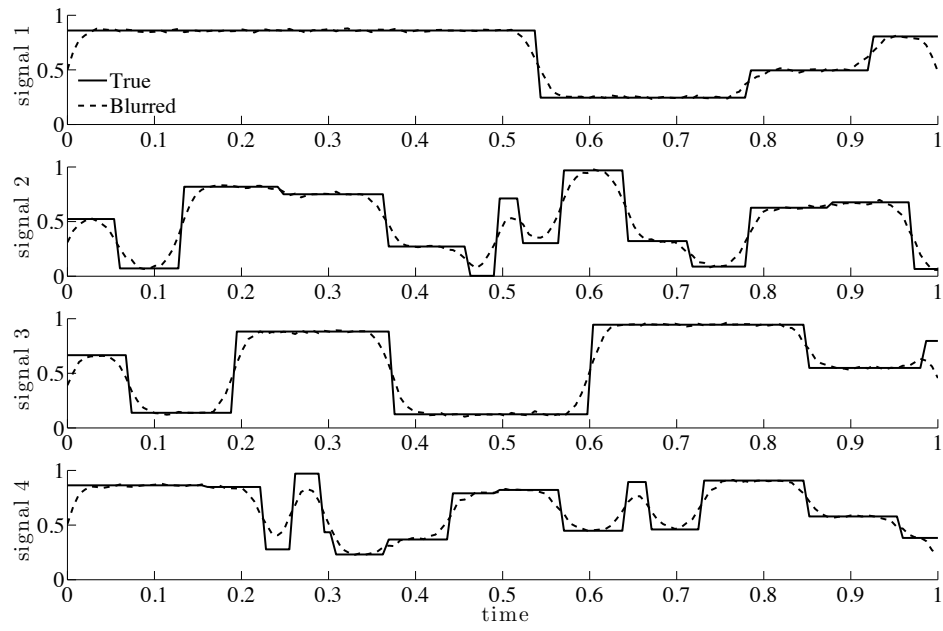


Figure 2: Four sample validation signals and corresponding observed signals, used in Experiments 1 and 2.

as the rank increases, but reconstructions degrade for larger ranks. On the other hand, since regularization is built into the design and construction of  $\tilde{\mathbf{Z}}_2$ , ORIM<sub>2</sub> avoids semi-convergence behavior altogether, as evident in the flattening of the error curve. Thus, ORIM<sub>2</sub> solutions are less sensitive to the choice of rank.

For ranks less than  $r = 57$ , TSVD- $\mathbf{A}$  achieves a smaller sample mean error than TSVD- $\hat{\mathbf{A}}$ . We also note that the smallest error that TSVD- $\hat{\mathbf{A}}$  can achieve is larger than the smallest error than TSVD- $\mathbf{A}$  can achieve. This is to be expected since  $\hat{\mathbf{A}}$  is an approximation of  $\mathbf{A}$ . Although the semi-convergence phenomenon for TSVD- $\hat{\mathbf{A}}$  occurs later and errors grow slower than for TSVD- $\mathbf{A}$ , both TSVD solutions require the selection of a good choice for the rank. Since training data is available, one could use the training data to select an “optimal” rank for TSVD- $\mathbf{A}$  and TSVD- $\hat{\mathbf{A}}$ , i.e., the rank which provides the minimal sample mean error for the training set. This is equivalent to the opt-TSVD approach from [14]. For TSVD- $\mathbf{A}$  and TSVD- $\hat{\mathbf{A}}$ , the minimal sample mean error for the training set occurs at ranks 51 and 54 (consistent with validation set), respectively.

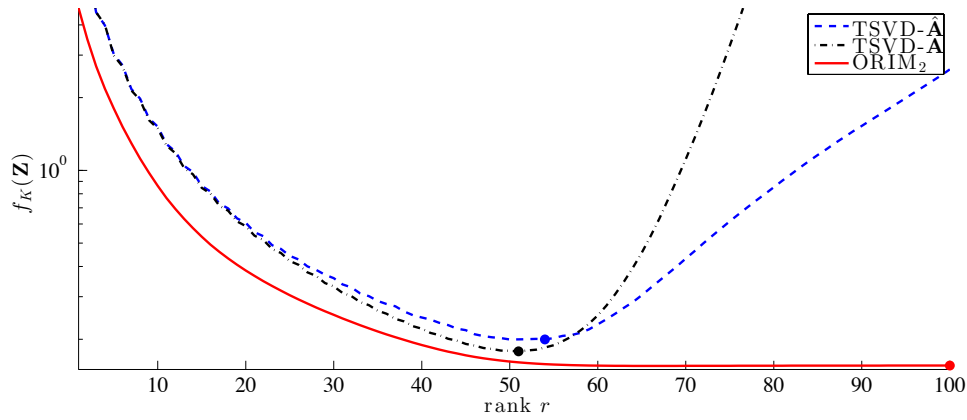


Figure 3: Sample mean error  $f_K$  on a logarithmic scale for ORIM<sub>2</sub>, TSVD- $\mathbf{A}$ , and TSVD- $\hat{\mathbf{A}}$  for the validation signals for various ranks. The dots correspond to the minimal sample mean error for the training data.

In Figure 4, we provide box-and-whisker plots to describe the statistics of the sample errors for the validation set,  $e_k(\mathbf{Z}) = \frac{1}{2} \left\| \mathbf{Z} \bar{\mathbf{b}}^{(k)} - \bar{\boldsymbol{\xi}}^{(k)} \right\|_2^2$  for  $k = 1, \dots, \bar{K}$ . Each box provides the median error, along with the 25<sup>th</sup> and 75<sup>th</sup> percentiles. The whiskers include extreme data, and outliers are plotted individually. For TSVD- $\mathbf{A}$  and TSVD- $\hat{\mathbf{A}}$ , errors correspond to the ranks computed using the opt-TSVD approach ( $r = 51$  and  $54$  respectively), and errors for ORIM<sub>2</sub> correspond to maximal rank 100. The average of the errors  $e_k(\mathbf{Z})$  correspond to the dots in Figure 3.

Next, we investigate how the errors for ORIM<sub>2</sub> vary for different numbers of training

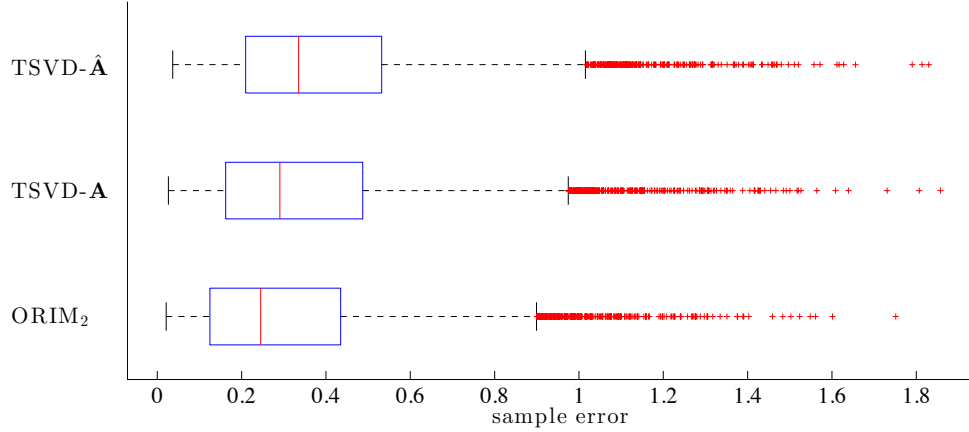


Figure 4: Box plots of sample errors for the validation set used in Experiment 1. All results use the 2-norm.

signals. That is, for increasing numbers of training signals, we compute the  $\text{ORIM}_2$  matrix and record sample errors  $e_k(\hat{\mathbf{Z}}_2)$ , for both the training set and the validation set. It is worthwhile to mention that the rank of  $\hat{\mathbf{Z}}_2$  is inherently limited by the number of training signals and the signal length, i.e.,  $r \leq \min\{K, n\}$ . For this example, we have  $n = 150$ , and we select the rank of  $\hat{\mathbf{Z}}_2$  to be

$$r = \begin{cases} K & \text{for } K < 50, \\ 50 & \text{for } K \geq 50. \end{cases} \quad (26)$$

Notice that for rank  $r = 50$  with a signal length of  $n = 150$ , we have a total of  $2 \cdot r \cdot n = 15,000$  unknowns (for  $\mathbf{X}$  and  $\mathbf{Y}$ ). In Figure 5 we provide the median errors along with the 25<sup>th</sup> to 75<sup>th</sup> percentiles. These plots are often referred to as Pareto curves in the literature.

We observe that for the number of training signals  $K < 50$ , the  $\text{ORIM}_2$  matrix is significantly biased towards the training set. This is evident in the numerically zero computed errors for the training set. In addition, we expect the errors for the training set to increase and the errors for the validation set to decrease as more training signals are included in the training set, due to a reduction in bias towards the training set. Eventually, errors for both signal sets reach a plateau, where additional training data do not significantly improve the solution. The point at which the errors begin to stabilize indicates a natural choice for the number of training signals to use, e.g., when the relative change in errors reaches a predetermined tolerance. For this example a relative improvement of  $10^{-3}$  for the errors for the training set is reached at 1,000 training images.

There is an interesting phenomena that occurs in the reconstruction errors for the validation set, whereby the error peaks around  $K = 150$ . Although puzzling at first, the peak in errors is intuitive when one considers the problem that is being solved. For  $K = 150$  matrix  $\mathbf{B}$  is  $150 \times 150$ . If  $\mathbf{B}$  is full rank ( $s = 150$ ),  $\mathbf{P} = \mathbf{C}$  in Theorem 3.1 and problem (19) has the

solution,

$$\begin{aligned}\hat{\mathbf{Z}}_2 &= \arg \min_{\text{rank}(\mathbf{Z}) \leq r} \frac{1}{K} \|\mathbf{Z}\mathbf{B} - \mathbf{C}\|_F^2 \\ &= \mathbf{C}_r \mathbf{B}^{-1},\end{aligned}$$

which is the special case considered by Eckart-Young and Schmidt [30, 31]. Notice that  $\hat{\mathbf{Z}}_2$  is uniquely determined by the training set, so reconstructions at this point will be most biased to the training set.

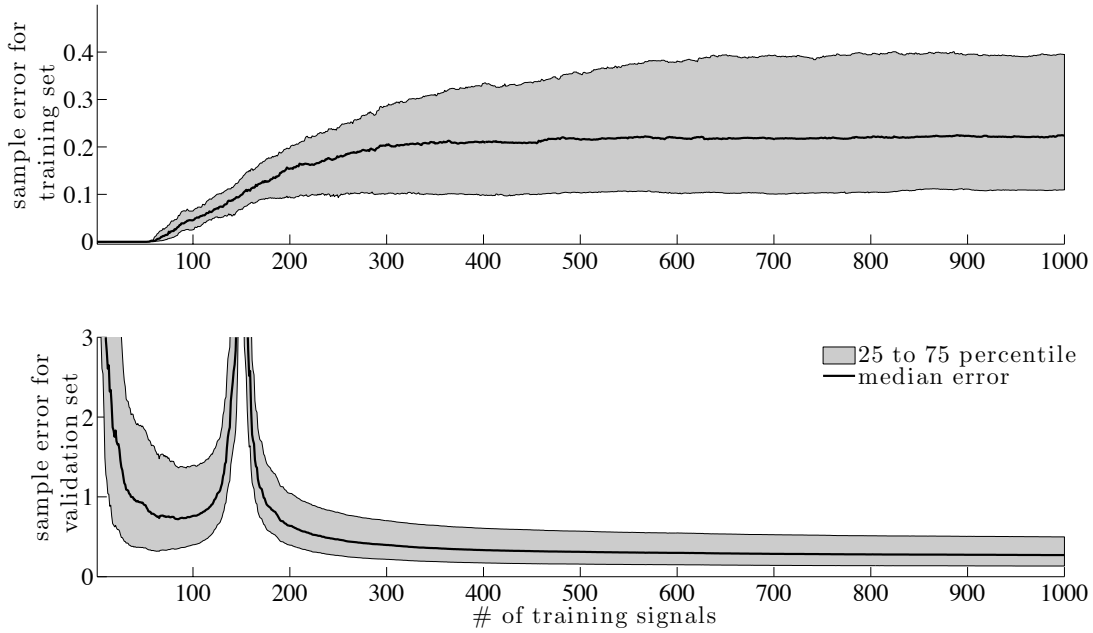


Figure 5: Pareto curves. The curves in the upper figure correspond to sample errors for the training set and the curves in lower figure correspond to sample errors for the validation set. Here we plot the median values along with the 25<sup>th</sup> and 75<sup>th</sup> percentiles of the errors as a function of the number of training signals used to construct  $\hat{\mathbf{Z}}_2$ .

**Experiment 2.** For this experiment, we use the same setup as described in Experiment 1 for the 1D example, but we compare results obtained by using different error measures  $\rho$  in equation (3). In particular, we consider  $\rho = \frac{1}{p} \|\cdot\|_p^p$  for  $p = 1.2, 2$ , and  $5$ . For  $p = 1.2$ , we use additional smoothing to deal with non-differentiability. We compute an optimal low-rank regularized inverse matrix  $\hat{\mathbf{Z}}_p$  for each  $p$ . The computed matrices are then used to reconstruct the validation set. We refer to these approaches as ORIM<sub>1.2</sub>, ORIM<sub>2</sub> and



ORIM<sub>5</sub>, respectively. Box-and-whisker plots of the sample errors for the validation signals are presented in Figure 6, along with errors for TSVD- $\hat{\mathbf{A}}$ . In each subfigure, the reconstruction error was computed using a different error measure. As expected, the reconstructions with smallest error correspond to  $\hat{\mathbf{Z}}_p$ , where  $p$  defines the error measure used for the training set.

We emphasize that even though the sample errors for the various methods may appear similar in value in the box plots, the reconstructions are significantly different, as seen in the absolute errors. Plots of the absolute values of the error vectors,  $\hat{\mathbf{Z}}_p \bar{\mathbf{b}}^{(k)} - \bar{\boldsymbol{\xi}}^{(k)}$ , for the four validation signals presented in Figure 2 can be found in Figure 7. We observe that ORIM<sub>5</sub> allows for large reconstruction errors throughout the signal without major peaks, while absolute reconstruction errors for ORIM<sub>1,2</sub> tend to remain small in flat regions, with only a few high peaks in the error at locations corresponding to discontinuities in the signal. These observations are consistent with norms in general and illustrate the potential benefits of using a different error measure.

## 4.2 Two-dimensional example

Next, we describe a problem from image deblurring, where the underlying model can be represented as (1) where matrix  $\mathbf{A}$  represents a blurring process, and  $\boldsymbol{\xi}$  and  $\mathbf{b}$  represent the true and observed images respectively. Entries in matrix  $\mathbf{A}$  depend on the point spread function (PSF). For image deblurring examples, matrix  $\mathbf{A}$  is typically very large and sparse. However, under assumptions of spatial invariance and assumed boundary conditions and/or properties of the PSF, structure in  $\mathbf{A}$  can be exploited and fast transform algorithms can be used to efficiently diagonalize  $\mathbf{A}$  [5]. For example, if the PSF satisfies a double symmetry condition and reflexive boundary conditions on the image are assumed, then  $\mathbf{A}$  can be diagonalized using the 2D discrete cosine transform (DCT). Of course, these approaches rely on knowledge of the exact PSF.

Blind deconvolution methods have been developed for problems where the PSF is not known or partially known [32]; however, most of these approaches require strict assumptions on the PSF such as spatial invariance or parameterization of the blur function, or they require a good initial estimate of the PSF. Our approach uses training data and does not require knowledge of  $\mathbf{A}$  and, hence, knowledge of the PSF. Furthermore, we do not require any assumptions regarding spatial invariance of the blur, boundary conditions, or special properties of the blur (e.g., separability, symmetry, etc.). We refer the reader to [24] for comparisons to blind deconvolution methods.

**Experiment 3.** In this experiment, we generate a set of 6,000 training images. Each image has  $128 \times 128$  pixels and represents a collection of circles, where the number of circles, location, radius and intensity were all randomly drawn from a uniform distribution. In this example, the PSF was a 2D Gaussian kernel with zero mean and variance 5. We assume

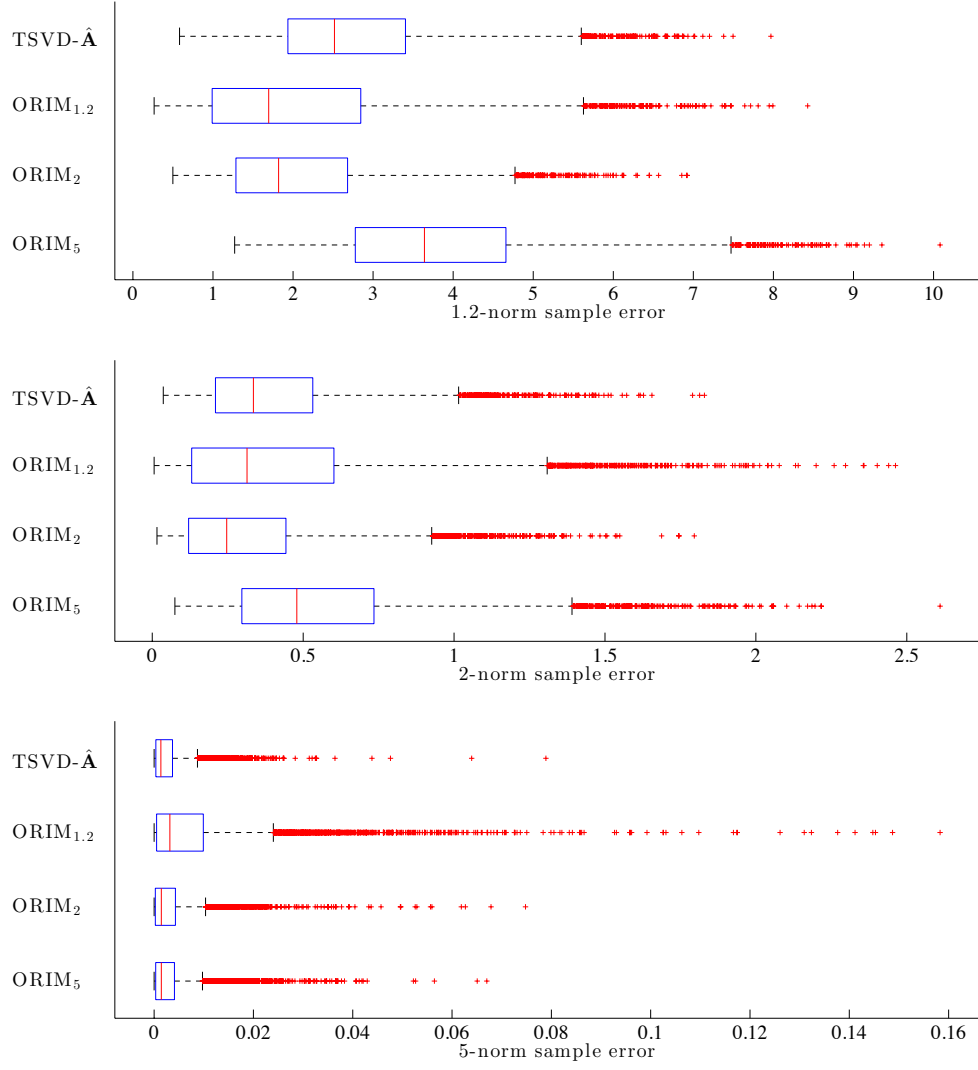


Figure 6: Box plots of sample errors for the validation set used in Experiment 2. Sample errors were computed using  $\rho = \frac{1}{p} \|\cdot\|_p^p$  where  $p = 1.2, 2$ , and  $5$ , respectively.

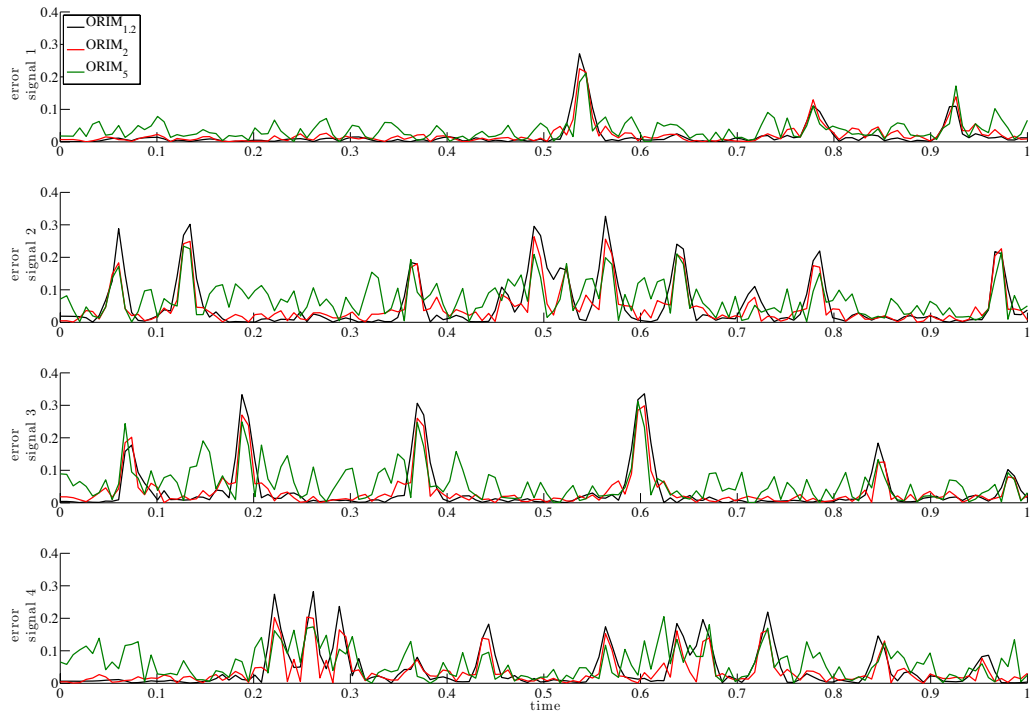


Figure 7: Absolute reconstruction errors for validation signals presented in Figure 2.

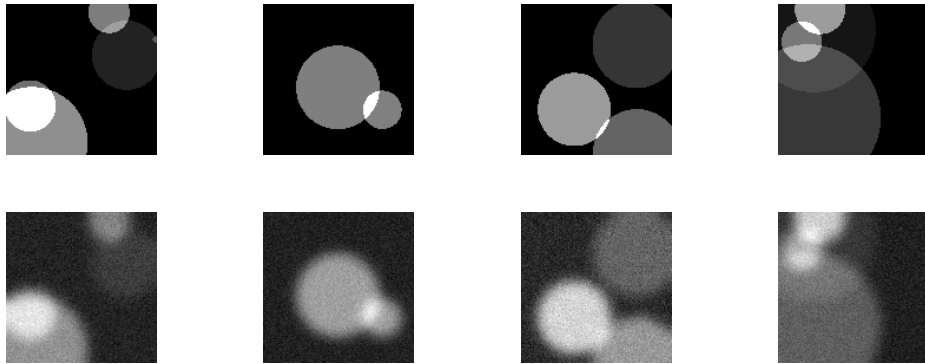


Figure 8: Sample true images and corresponding observed, blurred and noisy images that were used for training in Experiment 3.

reflexive boundary conditions for the image, and Gaussian white noise was added to each image, where the noise level was randomly selected between 0.1 and 0.15. Sample true images along with the corresponding blurred, noisy image can be found in Figure 8. A validation set of 7,000 images was generated in the same manner.

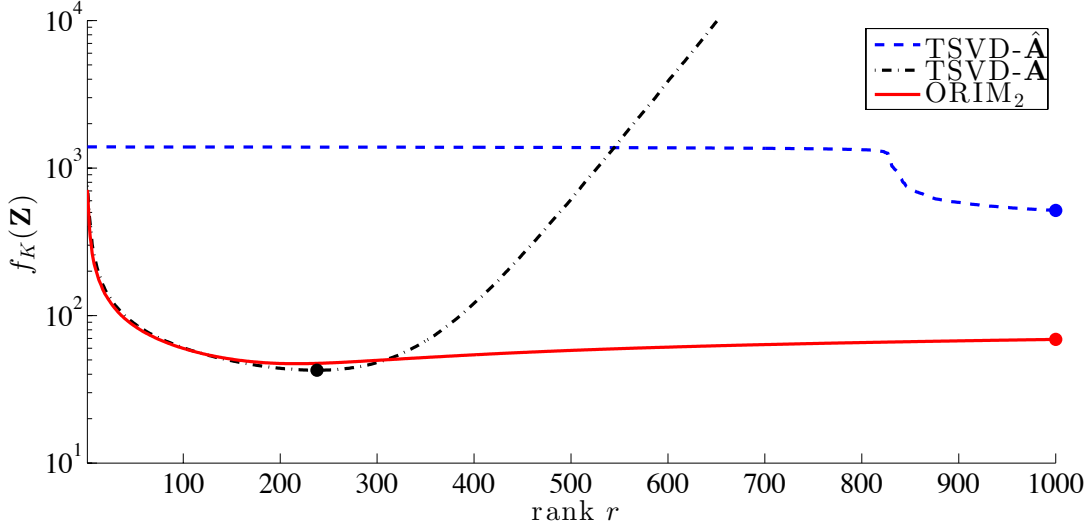


Figure 9: Sample mean error for  $\text{ORIM}_2$ ,  $\text{TSVD-}\mathbf{A}$ , and  $\text{TSVD-}\hat{\mathbf{A}}$  in logarithmic scale for validation data in Experiment 3. The dots correspond to the minimal sample mean error for the training data.

We compare the sample mean errors for the validation set for  $\text{ORIM}_2$ ,  $\text{TSVD-}\mathbf{A}$ , and  $\text{TSVD-}\hat{\mathbf{A}}$ , where the maximal rank for  $\hat{\mathbf{A}}$  was  $\bar{r} = 1,000$ . The sample mean errors are provided in Figure 9 for various ranks. Dots correspond to minimal sample mean error for the training data. As expected,  $\text{TSVD-}\mathbf{A}$  exhibits clear semi-convergence behavior, where the minimum error is achieved at rank 238 (for both the training and validation sets). In contrast,  $\text{ORIM}_2$  maintains low sample mean errors for various ranks, and only exhibits slight semiconvergence for larger ranks. This can be attributed to bias towards the training data, since for larger ranks, more training data is required to compensate for more unknowns. We observe that the relative difference between the lowest sample mean error for  $\text{ORIM}_2$  and the lowest sample mean error for  $\text{TSVD-}\mathbf{A}$  is approximately 10%. That is, even without knowledge of  $\mathbf{A}$ ,  $\text{ORIM}_2$  can produce reconstructions that are comparable to  $\text{TSVD-}\mathbf{A}$  reconstructions.

Next, we compare the distribution of the sample errors for  $\text{ORIM}_2$  and  $\text{TSVD-}\hat{\mathbf{A}}$ , corresponding to rank 1,000. Density plots can be found in Figure 10. Among the methods that only require training data,  $\text{ORIM}_2$  is less sensitive than  $\text{TSVD-}\hat{\mathbf{A}}$ , implying that it is better to seek a low rank regularized inverse matrix directly, versus seeking a low rank approximation to  $\mathbf{A}$  and then inverting it.

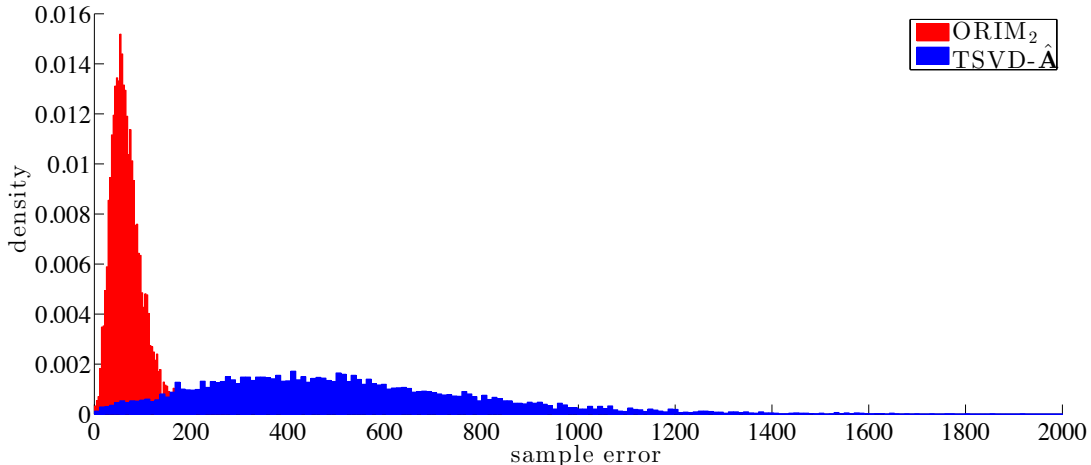


Figure 10: Density of sample errors for  $\text{ORIM}_2$  and  $\text{TSVD}-\hat{\mathbf{A}}$  for Experiment 3.

One of the validation images, with corresponding observed image can be found in the first column of Figure 11. Reconstructions for  $\text{ORIM}_2$ ,  $\text{TSVD}-\mathbf{A}$ , and  $\text{TSVD}-\hat{\mathbf{A}}$ , and corresponding absolute error images between the reconstructed image and the true image (with inverted colormap so that white corresponds to zero error) are provided in Figure 11 (top and bottom row respectively).

## 5 Conclusions

In this paper, we described a new framework for solving inverse problems where the forward model is not known, but training data is readily available. We considered the problem of finding an optimal low-rank regularized inverse matrix that, once computed, can be used to solve hundreds of inverse problems in real time. We used a Bayes risk minimization framework to provide underlying theory for the problem. For the empirical Bayes problem, we proposed a rank- $\ell$  update approach and showed that this approach is optimal in the Bayes framework for the 2-norm. Numerical results demonstrate that our approach can produce solutions which are superior to standard methods like TSVD, but without full knowledge of the forward model. The examples illustrate the benefits of using different error measures and the potential for use in large-scale imaging applications.

## Acknowledgment

We are grateful to Dianne P. O’Leary for helpful discussions.

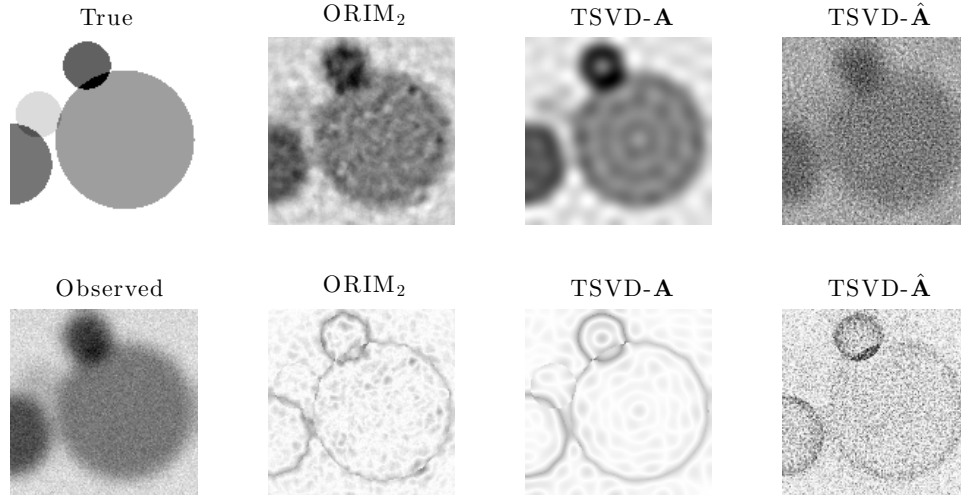


Figure 11: Image reconstructions for four validation images for Experiment 3. The bottom row presents absolute error images between the reconstructed image and the true image.

## References

- [1] H.W. Engl, M. Hanke, and A. Neubauer. *Regularization of Inverse Problems*. Springer, New York, 2000.
- [2] C.R. Vogel. *Computational Methods for Inverse Problems*. SIAM, Philadelphia, 2002.
- [3] A.G. Ramm. *Inverse Problems: Mathematical and Analytical Techniques with Applications To Engineering*. Springer, New York, 2004.
- [4] A. Tarantola. *Inverse Problem Theory and Methods for Model Parameter Estimation*. SIAM, Philadelphia, 2005.
- [5] P.C. Hansen, J.G. Nagy, and D.P. O’Leary. *Deblurring Images: Matrices, Spectra and Filtering*. SIAM, Philadelphia, 2006.
- [6] A. Kirsch. *An Introduction to the Mathematical Theory of Inverse Problems*. Springer, New York, 2011.
- [7] R.C. Aster, B. Borchers, and C.H. Thurber. *Parameter Estimation and Inverse Problems*. Academic Press, Waltham, 2012.
- [8] F.D.M. Neto and A.J.S. Neto. *An Introduction to Inverse Problems With Applications*. Springer, New York, 2012.

- [9] J. Hadamard. *Lectures on Cauchy's Problem in Linear Differential Equations*. Yale University Press, New Haven, 1923.
- [10] P.C. Hansen. *Rank-deficient and Discrete Ill-posed Problems*. SIAM, Philadelphia, 1998.
- [11] M. Mohri, A. Rostamizadeh, and A. Talwalkar. *Foundations of Machine Learning*. The MIT Press, Cambridge, MA, 2012.
- [12] Y.S. Abu-Mostafa, M. Magdon-Ismail, and H.T. Lin. *Learning from Data*. AMLBook, 2012.
- [13] V.N. Vapnik. *Statistical Learning Theory*. Wiley, San Francisco, 1998.
- [14] J. Chung, M. Chung, and D.P. O'Leary. Designing optimal filters for ill-posed inverse problems. *SIAM Journal on Scientific Computing*, 33(6):3132–3152, 2011.
- [15] J. Chung, M. Chung, and D.P. O'Leary. Optimal filters from calibration data for image deconvolution with data acquisition error. *Journal of Mathematical Imaging and Vision*, 44(3):366–374, 2012.
- [16] G.H. Golub and C.F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, 3 edition, 1996.
- [17] G.W. Stewart. *Matrix Algorithms: Volume 1, Basic Decompositions*. SIAM, Philadelphia, 1998.
- [18] B.P. Carlin and T.A. Louis. *Bayes and Empirical Bayes Methods for Data Analysis*. Chapman and Hall/CRC, Boca Raton, 2 edition, 2000.
- [19] A. Shapiro, D. Dentcheva, and R. Ruszczyński. *Lectures on Stochastic Programming: Modeling and Theory*. SIAM, Philadelphia, 2009.
- [20] G.A.F. Seber and A.J. Lee. *Linear Regression Analysis*. Wiley, San Francisco, 2 edition, 2003.
- [21] J. Chung, M. Chung, and D.P. O'Leary. Optimal regularized low rank inverse approximation. In revision at *Linear Algebra and Applications*, 2014.
- [22] D. Söndermann. Best approximate solutions to matrix equations under rank restrictions. *Statistische Hefte*, 27:57–66, 1986.
- [23] S. Friedland and A. Torokhti. Generalized rank-constrained matrix approximations. *SIAM Journal on Matrix Analysis and Applications*, 29(2):656–659, 2007.
- [24] J. Chung and M. Chung. Computing optimal low-rank matrix approximations for image processing. In *IEEE Proceedings of the Asilomar Conference on Signals, Systems, and Computers. November 3-6, 2013, Pacific Grove, CA, USA*, 2013.

- [25] J. Nocedal and S.J. Wright. *Numerical Optimization*. Springer, New York, 1999.
- [26] P.E. Gill, W. Murray, and M.H. Wright. *Practical Optimization*. Emerald Group Publishing, Bingley, UK, 1981.
- [27] P. Baldi and S. Brunak. *Bioinformatics: the Machine Learning Approach*. MIT press, Cambridge, 2001.
- [28] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning*. Springer, New York, 2 edition, 2009.
- [29] J.-F. Cai, E. J. Candès, and Z. Shen. A singular value thresholding algorithm for matrix completion. *SIAM Journal on Optimization*, 20(4):1956–1982, 2010.
- [30] C. Eckart and G. Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936.
- [31] E. Schmidt. Zur Theorie der linearen und nichtlinearen Integralgleichungen. *Mathematische Annalen*, 63(4):433–476, 1907.
- [32] P. Campisi and K. Egiazarian. *Blind Image Deconvolution: Theory and Applications*. CRC press, Boca Raton, 2007.